

We claim:

1. A system for storing checkpoint state information, comprising:
a network interface to an external network; and
a persistent memory unit coupled to the network interface, wherein:
the persistent memory unit is configured to receive the checkpoint data via a direct
memory write command from a primary process, and to provide access to
the checkpoint data via a direct memory read command from a backup
process, through the network interface; and
the backup process provides recovery capability in the event of a failure of the
primary process.
2. The system of Claim 1, further comprising:
a persistent memory manager configured to provide address context information to the
network interface.
3. The system of Claim 1, wherein the persistent memory unit is configured to
transmit the checkpoint data to another processor, and the backup process is executed by
the other processor.
4. The system of Claim 1, wherein the persistent memory unit provides the
checkpoint data upon request by the backup process when the primary process fails.
5. The system of Claim 1, wherein the persistent memory unit is configured to store
multiple sets of checkpoint data sent from the processor at successive time intervals.
6. The system of Claim 5, wherein the persistent memory unit provides the multiple
sets of checkpoint data upon request by the backup process at one time.
7. The system of Claim 1, wherein the primary process provides the checkpoint data
to the persistent memory unit independently from the backup process.

8. The system of Claim 1, wherein the persistent memory unit is configured as part of a remote direct memory access-enabled system area network.
9. The system of Claim 1, wherein the persistent memory unit is configured with address protection and translation tables to authenticate requests from remote processors, and to provide access information to authenticated remote processors.
10. A method for recovering the operational state of a primary process, comprising:
mapping virtual addresses of a persistent memory unit to physical addresses of the persistent memory unit;
receiving checkpoint data regarding the operational state of the primary process in the persistent memory unit; and
providing the checkpoint data to a backup process via a direct memory read command from the backup process.
11. The method of Claim 10, further comprising:
providing context information regarding the addresses to the primary process and the backup process.
12. The method of Claim 10, further comprising:
providing the checkpoint data to the backup process upon failure of the primary process.
13. The method of Claim 10, further comprising:
overwriting the checkpoint data with current checkpoint data.
14. The method of Claim 10, further comprising:
appending updated checkpoint data to at least one previous set of the checkpoint data.
15. The method of Claim 14, further comprising:
periodically supplying at least a portion of the multiple sets of checkpoint data in the backup process; and
clearing the portion of the multiple sets of checkpoint data.

16. The method of Claim 15, further comprising:
providing previously unread portions of the checkpoint data to the backup process upon
failure of the primary process; and
resuming functions performed by the primary process with the backup process.
17. The method of Claim 10, further comprising:
storing access information to the physical addresses of the checkpoint data in the
persistent memory unit when the primary process opens a memory region for the
checkpoint data; and
providing the access information to subsequent requestors of the checkpoint data.
18. The method of Claim 17, further comprising:
establishing a connection to a process requesting access to the checkpoint data; and
binding the access information to the connection.
19. The method of Claim 17, further comprising:
verifying authentication information from the subsequent requestors.
20. The method of Claim 10, further comprising:
authenticating a persistent memory manager during initialization of address protection
and translation tables on the persistent memory unit.
21. A computer product, comprising:
computer executable instructions operable to:
receive a direct memory access command from a remote processor via a network,
wherein the direct memory access command includes a reference to a
persistent memory virtual address;
receive checkpoint data from a primary process;
translate the virtual address to a physical address in the persistent memory unit;
and
allow access to the checkpoint data for use in a backup process.

22. The computer product of Claim 21, further comprising:
computer executable instructions operable to:
provide address context information to the processor.
23. The computer product of Claim 21, further comprising:
computer executable instructions operable to:
store multiple updates to the checkpoint data sent at successive time intervals.
24. The computer product of Claim 21, further comprising:
computer executable instructions operable to:
provides the multiple sets of checkpoint data to the backup process at one time.
25. The computer product of Claim 21, wherein the persistent memory is configured
as part of a remote direct memory access-enabled system area network.
26. An apparatus comprising:
means for communicatively coupling a persistent memory unit to a network that enables
direct access to the persistent memory unit;
means for mapping virtual addresses of the persistent memory unit to physical addresses
of the persistent memory unit;
means for receiving checkpoint data for a primary process in the persistent memory unit
via the network; and
means for providing the checkpoint data to a backup process via the network.
27. The apparatus of Claim 26, further comprising:
means for providing context information regarding the addresses to the primary process
and the backup process.
28. The apparatus of Claim 26, further comprising:
means for providing the checkpoint data to the backup process upon failure of the primary
process.

29. The apparatus of Claim 26, further comprising:
means for creating multiple sets of checkpoint data by appending updated checkpoint data to at least one previous set of the checkpoint data; and
means for overwriting the checkpoint data with current checkpoint data.
30. The apparatus of Claim 29, further comprising:
means for periodically supplying at least a portion of the multiple sets of checkpoint data in the backup process.
31. The apparatus of Claim 30, further comprising:
means for providing previously unread portions of the checkpoint data to the backup process upon failure of the primary process.
32. A method for recording the operational state of a primary process, comprising:
transmitting checkpoint data regarding the operational state of the primary process in the persistent memory unit via a direct memory access write command.
33. The method of Claim 32, further comprising:
overwriting the checkpoint data in the persistent memory unit with current checkpoint data via a direct memory access write command.
34. The method of Claim 32, further comprising:
appending updated checkpoint data to a previous set of the checkpoint data via a direct memory access write command.
35. A method for retrieving the operational state of a primary process, comprising:
transmitting a direct memory access read command via network to a remote persistent memory unit from a backup process for the primary process.
36. The method of Claim 35, further comprising:
periodically transmitting the direct memory access read command to retrieve at least a portion of the checkpoint data for the backup process.

37. The method of Claim 35, further comprising:
transmitting the direct memory access read command to retrieve previously unread
portions of the checkpoint data upon failure of the primary process.